



US007133857B1

(12) **United States Patent**
Zhang et al.

(10) **Patent No.:** US 7,133,857 B1
(45) **Date of Patent:** Nov. 7, 2006

(54) **PROCESSING STRUCTURED DATA**(75) Inventors: **Zhengyu Zhang**, Milpitas, CA (US);
Hui Tian, Milpitas, CA (US)(73) Assignee: **Ximpleware, Inc.**, Milpitas, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 477 days.

(21) Appl. No.: **10/272,077**(22) Filed: **Oct. 15, 2002**(51) **Int. Cl.****G06F 7/08** (2006.01)(52) **U.S. Cl.** **707/1; 707/100; 715/513**(58) **Field of Classification Search** **707/1; 707/100; 715/513**
See application file for complete search history.(56) **References Cited**

U.S. PATENT DOCUMENTS

- 6,209,124 B1 * 3/2001 Vermeire et al. 717/114
2002/0029229 A1 * 3/2002 Jakopac et al. 707/500
2002/0038319 A1 * 3/2002 Yahagi 707/513
2002/0087596 A1 * 7/2002 Lewontin 707/513

2002/0143521 A1 * 10/2002 Call 704/1

OTHER PUBLICATIONS

H. Blume et al. Integration of High-Performance ASICs into Reconfigurable Systems Providing Additional Multimedia Functionality, 2000, IEEE, 10 pages.*

M. Morris Mano, Computer System Architecture, Prentice-Hall, Third Edition, 4 pages, 1993.*

* cited by examiner


Primary Examiner—Paul Rodriguez

Assistant Examiner—Luke Osborne

(74) Attorney, Agent, or Firm—Beyer Weaver & Thomas, LLP

(57) **ABSTRACT**

The present invention provides a fast and efficient way of processing structured data by utilizing an intermediate file to store the structural information. The structured data may be processed into a Binary mask Format (BMF) file which may serve as a starting point for post-processing. A tree structure built on top of the BMF file may be constructed very quickly, and also takes up less space than a *Document Object Model* (DOM) tree. Additionally, BMF records may reside entirely in the memory and contain structural information, allowing *Simple API for XML* (SAX)-like sequential data access.

43 Claims, 6 Drawing Sheets

U.S. Patent

Nov. 7, 2006

Sheet 1 of 6

US 7,133,857 B1




FIG. 1

U.S. Patent

Nov. 7, 2006

Sheet 2 of 6

US 7,133,857 B1




FIG. 2

U.S. Patent

Nov. 7, 2006

Sheet 3 of 6

US 7,133,857 B1




FIG. 3

U.S. Patent

Nov. 7, 2006

Sheet 4 of 6

US 7,133,857 B1





FIG. 4

U.S. Patent

Nov. 7, 2006

Sheet 5 of 6

US 7,133,857 B1**FIG. 5**

U.S. Patent

Nov. 7, 2006

Sheet 6 of 6

US 7,133,857 B1




FIG. 6

US 7,133,857 B1

1

PROCESSING STRUCTURED DATA

FIELD OF THE INVENTION

The present invention relates to the field of structured data files in computer systems. More specifically, the present invention relates to the processing of structured data in an efficient manner.

BACKGROUND OF THE INVENTION

Structured data represents a large portion of the information accessed on the Internet and other computer networks. There are several reasons why structured data is so popular. American Standard Code for Information Interchange (ASCII) and its extensions, such as Unicode Transformation Formats UTF-8 and UTF-16 are among the most common standard encoding formats. Text encoding puts information into a format that is easily readable by a human, thus it is easy for programmers to develop and be applications. Lastly, textual encoding is extensible and adding new information may be as simple as adding a new key-value pair.

Recently, Extensible Markup Language (XML) has been growing in popularity. XML is a markup language for documents containing structured information. Unlike its predecessor, Hypertext Markup Language (HTML), where tags are used to instruct a web browser how to render data, in XML the tags are designed to describe the data fields themselves. XML, therefore, provides a facility to define tags and the structural relationships between them. This allows a great deal of flexibility in defining markup languages to using information. Because XML is not designed to do anything other than describe what the data is, it serves as the perfect data interchange format.

XML, however, is not without its drawbacks. Compared with other data formats, XML can be very verbose. Processing an XML file can be very CPU and memory intensive, severely degrading overall application performance. Additionally, XML suffers many of the same problems that other software-based text-based processing methods have. Modern processors prefer binary data representations, particularly ones that fit the width of the registers, over text-based representations. Furthermore, the architecture of many general-purpose processors trades performance for programmability, thus making them ill-suited for text processing. Lastly, the efficient parsing of structured text, no matter the format, can present a challenge because of the added steps required to handle the structural elements.

Most current XML parsers are software-based solutions that follow either the Document Object Model (DOM) or Simple API for XML (SAX) technologies. DOM parsers convert an XML document into an in-memory hierarchical representation (known as a DOM tree), which can later be accessed and manipulated by programmers through a standard interface. SAX parsers, on the other hand, treat an XML document as a stream of characters. SAX is event-driven, meaning that the programmer specifies an event that may happen, and if that event occurs, SAX gets control and handles the situation.

In general, DOM and SAX are complementary, not competing, XML processing models, each with its own benefits and drawbacks. DOM programming is programmer-friendly, as the processing phase is separate from application logic. Additionally, because the data resides in the memory, repetitive access is fast and flexible. However, DOM requires that the entire document data structure, usually

2

occupying 7–10 times the size of the original XML document, be loaded into the memory, thus making it impractical for large XML documents. SAX, on the other hand, can be efficient in parsing large XML documents (at least when only small amounts of information need to be processed at once), but it maintains little of the structural information of the XML data, putting more of a burden on programmers and resulting in code that is hardwired, bulky, and difficult to maintain.

What is needed is an application program interface (API) that combines the best attributes of both DOM and SAX parsing.

BRIEF DESCRIPTION OF THE INVENTION

The present invention provides a fast and efficient way of processing structured data by utilizing an intermediate file to store the structural information. The structured data may be processed into a Binary mask Format (BMF) file which may serve as a starting point for post-processing. A tree structure built on top of the BMF file may be constructed very quickly, and also takes up less space than a DOM tree. Additionally, BMF records may reside entirely in the memory and contain structural information, allowing SAX-like sequential data access.

BRIEF DESCRIPTION OF THE DRAWINGS

The accompanying drawings, which are incorporated into and constitute a part of this specification, illustrate one or more embodiments of the present invention and, together with the detailed description, serve to explain the principles and implementations of the invention.

In the drawings:

FIG. 1 is a block diagram illustrating a layer view of an apparatus for efficiently processing structured data in accordance with an embodiment of the present invention.

FIG. 2 is a block diagram illustrating an apparatus for efficiently processing structured data in accordance with an embodiment of the present invention.

FIG. 3 is a timing diagram illustrating the operation of the hardware in accordance with an embodiment of the present invention.

FIG. 4 is a flow diagram illustrating a method for modifying the content of a target string in a BMF file from an old string to a new string in accordance with an embodiment of the present invention.

FIG. 5 is a diagram illustrating a BMF record format in accordance with an embodiment of the present invention.

FIG. 6 is a flow diagram illustrating a method for efficiently processing a structured data file, the structured data file including one or more pieces of content, in accordance with an embodiment of the present invention.

DETAILED DESCRIPTION

Embodiments of the present invention are described herein in the context of a system of computers, servers, and software. Those of ordinary skill in the art will realize that the following detailed description of the present invention is illustrative only and is not intended to be in any way limiting. Other embodiments of the present invention will

US 7,133,857 B1

3

readily suggest themselves to such skilled persons having the benefit of this disclosure. Reference will now be made in detail to implementations of the present invention as illustrated in the accompanying drawings. The same reference indicators will be used throughout the drawings and the following detailed description to refer to the same or like parts.

In the interest of clarity, not all of the routine features of the implementations described herein are shown and described. It will, of course, be appreciated that in the development of any such actual implementation, numerous implementation-specific decisions must be made in order to achieve the developer's specific goals, such as compliance with application- and business-related constraints, and that these specific goals will vary from one implementation to another and from one developer to another. Moreover, it will be appreciated that such a development effort might be complex and time-consuming, but would nevertheless be a routine undertaking of engineering for those of ordinary skill in the art having the benefit of this disclosure.

In accordance with the present invention, the components, process steps, and/or data structures may be implemented using various types of operating systems, computing platforms, computer programs, and/or general purpose machines. In addition, those of ordinary skill in the art will recognize that devices of a less general purpose nature, such as hardwired devices, field programmable gate arrays (FPGAs), application specific integrated circuits (ASICs), or the like, may also be used without departing from the scope and spirit of the inventive concepts disclosed herein.

For purposes of this disclosure, a structured data file is any file containing content as well as some information regarding the structural organization of the content. The present invention provides a fast and efficient way of processing structured data by utilizing an intermediate file to store the structural information. The structured data may be processed into a Binary mask Format (BMF) file which may serve as a starting point for post-processing. A tree structure built on top of the BMF file may be constructed very quickly, and also takes up less space than a DOM tree. Additionally, BMF records may reside entirely in the memory and contain structural information, allowing SAX-like sequential data access. However, while this document will describe advantages that the present invention provides over DOM or SAX, one of ordinary skill in the art will recognize that the present invention need not be limited to replacing DOM or SAX, and can be expanded to non-XML type processing.

FIG. 1 is a block diagram illustrating a layer view of an apparatus for efficiently processing structured data in accordance with an embodiment of the present invention. The apparatus may comprise three layers. A hardware text processing accelerator 100, occupying the lowest layer, may offer the horsepower necessary to relieve the central processing unit (CPU) from the most processor intensive part of the task. On top of the hardware text processing accelerator 100 may lie a device driver layer 102 that is responsible for the communication between the hardware text processing accelerator 100 and a software layer 104. The software layer 104 may be designed to offer maximum flexibility and further improve the performance. It may export APIs that are standard-compliant.

The hardware may be designed such that it may quickly match multiple patterns against an incoming data stream. FIG. 2 is a block diagram illustrating an apparatus for efficiently processing structured data in accordance with an embodiment of the present invention. A text processor 200 may be the core of the accelerator. It may be composed of

4

multiple Finite State Machines (FSMs) that process an incoming document in parallel. The output may be the BMF file. It also may set several result registers (not pictured). A PCI interface 202 may handle all handshaking between the hardware and a server PCI bus 204. A memory controller 206 may receive commands from the CPI interface 202 and convert the PCI address to on-board memory address space. It also may access the board memory accordingly. Configuration registers 208 may determine the configuration of the test processing pipeline and the organization of the memory controller 206. It may load default values from configuration ROM 210. Some of the values may be modified by software through the PCI interface 202. The Configuration ROM 210 may store the default setting of the text processor configuration. It also may store the configuration map of FPGAs.

A document buffer 212 may store the incoming document. This may be a Synchronous Dynamic Random Access Memory (SDRAM). Paging may be utilized if the incoming document is larger than the total buffer size. A BMF buffer 214 may store the output BMF files, together with several other text processor result register values. This may be a separate SDRAM, although in one embodiment of the present invention it may share a single SDRAM with the document buffer. A string cache 216 may buffer the incoming data to smooth out and speed up SDRAM access. A DMA engine 218 may grab the frame data from server main memory and send it back the BMF file.

FIG. 3 is a timing diagram illustrating the operation of the hardware in accordance with an embodiment of the present invention. A reset 300 may be sent out by the host computer system, or by a specific application. While reset is asserted, the configuration may be loaded from ROM to the configuration registers. The text processor then may be set to a default state according to the configuration registers. When software calls the device driver, a start signal 302 may be sent through the PCI to indicate the beginning of a document processing cycle. Then the PCI master may assert a frame number 304 to indicate the beginning of a write transaction. The PCI master may then drive the address/data 306 to the PCI bus. The PCI target interface may respond, causing the DMA to read the document into the SDRAM document buffer 308. There may also be certain PCI commands reserved to update the configuration registers. The memory controller may be activated by PCI command. It may start processing data in the string buffer 310. It also may send sync signals to the memory controller. The transferring and processing may be repeated.

The PCI target may sense a valid window to send data. Then the PCI master may assert the frame number to indicate the beginning of a read transaction 312. The PCI target holds the bus. The DMA engine may then transfer the BMF and result register data to main memory 314. When all the data is transferred, the PCI target interface may send an end signal to the device driver 316. The next document processing cycle may start again with a start signal from the device driver 318.

The output of the hardware is a BMF. In one embodiment of the present invention, the BMF defines a binary record format that is used to describe various fields in a structured data file. It can be viewed as a two-dimensional field of bits. Each piece of useful information in the structured data file may correspond to a record in the BMF file. A record may comprise a starting offset and length of a target string in the structured data file. It may also comprise the depth value,

US 7,133,857 B1

5

node type, and bit-wide flanges. These will be discussed in more detail below. The total length of a record may be an integer multiple of a 32-bit word—the width of the memory bus in most commercial architectures. Two record types may be defined: a full version of 12 bytes in length, and a compressed version of 8 bytes in length. The full version may be based on the assumption that both the string length and the starting offset are 32-bits wide, whereas the compressed version may assume a 16-bit field, which translates to a maximum length of 65536 bytes.

FIG. 4 is a flow diagram illustrating a method for modifying the content of a target string in a BMF file from an old string to a new string to a new string in accordance with an embodiment of the present invention. At 400, a piece of memory of the length of the new string may be allocated. At 402, the new string may be filled into the memory. At 404, a record in the BMF file corresponding to the old string may be located. At 406, a corresponding bit flag for the record may be changed from relative offset to absolute offset. The first time a file is converted to BMF form, relative offsets may be used. However, as soon as a modification is made to a specific string, it can often be difficult if not impossible to continue to track the relative offset for that string. Therefore, absolute offsets may be utilized for all modified strings. At 408, an offset value in the BMF record for the old string may be replaced with a pointer value for the new string in memory. At 410, a length field in the VMF record for the old string may be replaced with the length of the new string.

There are at least three types of possible BMF file modes: read-only mode, read-modify mode, and read-modify-add mode. In read-only mode, records representing various types of nodes may be placed sequentially into a BMF file, leaving no empty records. For example, a leaf-level element may be represented as a record for the starting tag, 2n records (one of r property name and one property value) for n properties, and one record for the text for the element, and finally one record for an ending tag name. The presence of the ending tag record may be used for document validation.

The read-modify mode may be identical to read-only mode except each record allows for limited write-access, meaning content can be altered, but not added.

The read-modify-add mode allows complete write-access, which is done by embedding empty records into the file.

The record format may be picked to efficiently represent the necessary information of the original data structure. It may also be made less efficient on purpose to speed up downstream processing.

FIG. 5 is a diagram illustrating a BMF record format in accordance with an embodiment of the present invention. A control word may be thirty-two bits in total in this embodiment. A starting offset 500 may indicate a starting position relative to the beginning of a structured data file. Additionally, the control word may include a depth value 502 of sixteen bits. The depth value may indicate the depth of a tag in the hierarchy of tags in the structured data file. Thus, the first tag in a file will have a depth of zero, whereas if another starting tag appears before an ending tag for the first tag, that second starting tag will have a depth of one. A content type 504 may be provided, which indicates what type of information the content is. In an embodiment of the present invention, the following value/content type pairs may be used in this field:

6

TABLE 1

<u>Content Types and Corresponding Values</u>				
Content	Type	Value	Content name	Example
10	Starting Tag	0	<example>	<example>
	Ending Tag	1	</example>	</example>
	Property Name	2	<example property1="this">	<example property1="this">
	Property Value	3	<example property2="that">	<example property2="that">
	Text	4	<example> tastyfruit </example>	<example> tastyfruit </example>
	Comment	5	<!-- - thisisacomment -->	<!-- - thisisacomment -->
	Processing Instruction	6	<??>	<??>
	Markup declaration I name	7	<![[CDATA[...<<>>...]]>	<![[CDATA[...<<>>...]]>
	Markup declaration I value	8	<![[CDATA[...<<>>...]]>	<![[CDATA[...<<>>...]]>
	Markup declaration II name	9	<!ENTITY...>	<!ENTITY...>
	Markup declaration II value	10	<!ENTITY ...>	<!ENTITY ...>
	Entity reference	11	&example.bib;	&example.bib;
	Property Name Value Pair	12	<example property1="this">	<example property1="this">

20 A modification indicator 506 may also be provided, which indicates whether or not the record has been modified. This is important because, as described above, if the record has been modified, then the offset field will contain the real pointer value, not a relative offset. An insertion indicator 508 may indicate that data was inserted in between two existing records. Once again, this is important in determining how to utilize the offset field. If the insertion indicator is set to 1, it indicates that the offset field contains a pointer to an external piece of memory, one that can be used to add child nodes to the current node. At end of document indicator 510 may indicate whether the tag is the last one in the document. This can be important because in some embodiments, ending tags may be ignored when encoding the BMF file in order to save space. Therefore, the last tag in the BMF file may not correspond to the last tag in the structured data file.

25 A current record in use field 512 may be used to indicate that a record has been deleted. If the field is set to 0, the record may be safely ignored because it has been deleted. A name space indicator 514 may indicate whether or not there is a name space within the token (which may be represented by an “=” sign). A reference bit 516 may indicate when there is an external reference, such as an “&” in a text string.

30 There may be one or more reserved bits 518, which are set aside for future uses. Lastly, a length field 520 may indicate the length of the content.

35 The BMF file, together with the original data in memory, completely describes the original data and its inherent data structure. Traversing the data structure may be easily accomplished using the BMF records. Higher level applications and processing are therefore facilitated by using the BMF. To make it easily accessible and readily integrated to higher level application and processing, device drivers and an application programming interface (API) may be built on top of the BMF.

40 FIG. 6 is a flow diagram illustrating a method for efficiently processing a structured data file, the structured data file including one or more pieces of content, in accordance with an embodiment of the present invention. The structured data file may be an extensible markup language file. The process loops through each piece of content. In another embodiment of the present invention, the process loops through each relevant piece of content. Relevancy can be determined by the programmer and may be chosen so as to minimize the amount of space used for a BMF file. At 600, a BMF record is created in a BMF file, the BMF record corresponding to the piece of content. At 602, an offset may

US 7,133,857 B1

7

be stored in the BMF record indicating a starting position for the piece of content relative to the beginning of the structured data file. At 604, a depth of the piece of content may be stored in the BMF record, the depth indicating a level in a hierarchy of tags in the structured data file. At 606, a content type of the piece of content may be stored in the BMF record, the content type indicating a type of information for the piece of content. The content type may take many forms, such as a starting tag, ending tag, property name, property value, text, comment, processing instruction, markup declaration name, markup declaration value, external reference, property name pair, etc. At 608, a length may be stored for the piece of content in the BMF record.

At 610, a modification indicator for the piece of content may be stored in the BMF record, the modification indicator indicating if the BMF record has been modified and the modification indicator initially set to indicate that no modification has been made. At 612, an insertion indicator for the piece of content may be stored in the BMF record, the insertion indicator indicating if the BMF record has been

5

10

15

20

8

The following example may be used to illustrate an embodiment of the present invention. One of ordinary skill in the art will recognize that this is merely an example and should not be read to be limiting in any way. Suppose an XML file as follows:

```
<?xml version="1.0" encoding="US-ASCII"?>
<benchmark:database xmlns:benchmark="http://example.com/xml/benchmark">
  benchmark:person id="012345">
    benchmark:email ame="Name012345"/>
  <!--Edited with XML spy v4.2-->
  <benchmark:line1>Line 1012345 012345</benchmark:line1>
</benchmark:person>
</benchmark:database>
```

An embodiment of the present invention may ignore ending tags and produce the following BMF file:

Starting offset 32 bit	Depth 5 bit	Token type 4 bit	Modify indicator 1 bit	Insertion Indicator 1 bit	End of Document Indicator 1 bit	Current Record in use 1 bit	Name space indicator 1 bit	Reference indicator 1 bit	unused	Length 16 bit
2	0	6	0	0	0	1	0	0	0	38
42	0	0	0	0	0	1	1	0	0	18
61	0	2	0	0	0	1	1	0	0	14
78	0	3	0	0	0	1	0	0	0	35
116	1	0	0	0	0	1	1	0	0	16
133	1	2	0	0	0	1	0	0	0	2
137	1	3	0	0	0	1	0	0	0	6
147	2	0	0	0	0	1	1	0	0	15
163	2	2	0	0	0	1	0	0	0	4
169	2	3	0	0	0	1	0	0	0	10
185	1	5	0	0	0	1	0	0	0	25
218	2	0	0	0	0	1	1	0	0	15
234	2	4	0	0	0	1	0	0	0	23
0	0	0	0	0	1	1	0	0	0	0

The packet BMF records are:

0000000000000000000000000000000010	00000	0110	0	0	0	1	0	0	0	00100110
00000000000000000000000000000000101010	00000	0000	0	0	0	1	1	0	0	00010010
00000000000000000000000000000000111101	00000	0010	0	0	0	1	1	0	0	00001110
000000000000000000000000000000001001110	00001	0011	0	0	0	1	0	0	0	00100011
000000000000000000000000000000001110100	00001	0000	0	0	0	1	1	0	0	00010000
0000000000000000000000000000000010000101	00001	0010	0	0	0	1	0	0	0	00000010
0000000000000000000000000000000010001001	00010	0011	0	0	0	1	0	0	0	00000110
0000000000000000000000000000000010010011	00010	0000	0	0	0	1	1	0	0	00001111
0000000000000000000000000000000010100011	00010	0010	0	0	0	1	0	0	0	00000100
0000000000000000000000000000000010101001	00001	0011	0	0	0	1	0	0	0	000001010
0000000000000000000000000000000010111001	00010	0101	0	0	0	1	0	0	0	00011001
0000000000000000000000000000000011011010	00010	0000	0	0	0	1	1	0	0	00001111
0000000000000000000000000000000011101010	00010	0100	0	0	0	1	0	0	0	00010111
00	00000	0000	0	1	1	1	0	0	0	00000000

inserted between two existing BMF records and the insertion indicator initially set to indicate that the BMF record has not been inserted between two existing BMF records. At 614, an end of document indicator for the piece of content may be stored in the BMF record, the end of document indicator indicating if the BMF record corresponds to a last piece of content in the structured data file. At 616, a current record in use field may be stored for the piece of content in the BMF record, the current record in use field indicating whether the piece of content has been deleted.

While embodiments and applications of this invention have been shown and described, it would be apparent to those skilled in the art having the benefit of this disclosure that many more modifications than mentioned above are possible without departing from the inventive concepts herein. The invention, therefore, is not to be restricted except in the spirit of the appended claims.

What is claimed is:

1. A method for efficiently processing a structured data file, the structured data file including one or more pieces of content, comprising:

US 7,133,857 B1

9

creating a record in an intermediate file for a piece of content in the structured data file, the record containing an offset indicating a starting position for the piece of content in the structured data file and a depth value of the piece of content indicating a level in a hierarchy of tags in the structured data file; and

formatting said record in a way that allows data to be accessed using both said intermediate file and the structured data file without traversing the entire structured data file to determine the depth value.

2. The method of claim 1, wherein said creating a record comprises; creating a binary mask format (BMF) record in the intermediate file, the BMF record corresponding to said piece of content; and

storing an offset in said BMF record indicating a position for said piece of content relative to a point in said structured data file.

3. The method of claim 1, wherein said intermediate file is a BMF file, wherein the BMF file includes one or more BMF records.

4. The method of claim 2, wherein said creating a record further comprises: storing a content type of said piece of content in said BMF record, said content type indicating a type of information for said piece of content.

5. The method of claim 2, wherein said creating a record further comprises: storing a length for said piece of content in said BMF record.

6. The method of claim 2, wherein said offset indicates a starting position for said piece of content relative to a beginning of the structured data file.

7. A method for efficiently processing a structured data file, the structured data file including one or more pieces of content, comprising:

creating a binary mask formed (BMF) record in a BMF file, said BMF record corresponding to a piece of content;

storing an offset in said BMF record indicating a starting position for said piece of content relative to a beginning of the structured data file;

storing a depth value of said piece of content in said BMF record, said depth value indicating a level in a hierarchy of tags in said structured data file;

storing a content type of said piece of content in said BMF record, said content type indicating a type of information for said piece of content; and

storing a length for said piece of content in said BMF record; and

wherein said BMF record may be accessed to determine the depth value without traversing the entire structured data file.

8. The method of claim 7, further comprising:

storing a modification indicator for said piece of content in said BMF record, said modification indicator indicating if said BMF record has been modified and said modification indicator initially set to indicate that no modification has been made.

9. The method of claim 7, further comprising:

storing an insertion indicator for said piece of content in said BMF record, said insertion indicator indicating if said BMF record has been inserted between two existing BMF records and said insertion indicator initially set to indicate that said BMF record has not been inserted between two existing BMF records.

10. The method of claim 7, further comprising:

storing an end of document indicator for said piece of content in said BMF record, said end of document

10

indicator indicating if said BMF record corresponds to a last piece of content in said structured data file.

11. The method of claim 7, further comprising: storing a current record in use field for said piece of content in said BMF record, said current record in use field indicating whether said piece of content has been deleted.

12. The method of claim 7, wherein said content type is a tag, a property, text, or a comment.

13. The method of claim 7, wherein said content type is a starting tag, property name, property value, text, comment, or processing instruction.

14. The method of claim 7, wherein said content type is a starting tag, ending tag, property name, property value, text, comment, processing instruction, markup declaration name, markup declaration value, external reference, or property name value pair.

15. The method of claim 7, wherein said structured data file is an extensible markup language file.

16. An apparatus for efficiently processing structured data, comprising:

a peripheral component interface (PCI) interface; a direct memory access (DMA) engine coupled to said PCI interface;

a text processor coupled to said PCI interface, the text processor configured to:

create a record in an intermediate file for a piece of content in the structured data file, the record containing an offset indicating a starting position for the piece of content in the structured data file and a depth value of the piece of content indicating a level in a hierarchy of tags in the structured data file; and

format said record in a way that allows data to be accessed using both said intermediate file and the structured data file without traversing the entire structured data file to determine the depth value;

configuration memory coupled to said text processor and to said PCI interface;

a memory controller coupled to said PCI interface; BMF memory coupled to said DMA engine, said memory controller, and said text processor;

a document buffer coupled to said DMA engine, said memory controller, and said text processor; and a string cache coupled to said DMA engine, said memory controller, and said text processor.

17. The apparatus of claim 16, wherein said configuration memory includes: one or more configuration registers; and configuration read-only-memory coupled to said one or more configuration registers.

18. The apparatus of claim 16, wherein said PCI interface is configured to handle all handshaking between the apparatus and a server PCI bus.

19. The apparatus of claim 16, wherein said memory controller is configured to receive commands from said PCI interface and convert a PCI address to on-board memory address space.

20. The apparatus of claim 19, wherein said memory controller is further configured to access board memory according to said PCI address.

21. The apparatus of claim 19, wherein said configuration register contains a configuration of a test processing pipeline and organization of said memory controller.

22. The apparatus of claim 19, wherein said memory buffer is configured to store an incoming document.

23. The apparatus of claim 19, wherein said memory buffer is Synchronous Dynamic Random Access memory (SDRAM).

US 7,133,857 B1

11

24. The apparatus of claim **19**, wherein said DMA engine is configured to grab frame data from server main memory and send back a BMF file.

25. The apparatus of claim **19**, wherein said text processor is configured to, for each relevant piece of content in the structured data file:

create a binary mask format (BMF) record in a BMF file, said BMF record corresponding to said piece of content;

store an offset in said BMF record indicating a starting position for said piece of content relative to a beginning of the structured data file;

store a depth of said piece of content in said BMF record, said depth indicating a level in a hierarchy of tags in said structured data file;

store a content type of said piece of content in said BMF record, said content type indicating a type of information for said piece of content; and

store a length for said piece of content in said BMF record.

26. An apparatus for efficiently processing a structured data file, the structured data file including one or more pieces of content, comprising:

means for creating a record in an intermediate file for a piece of content in the structured data file, the record containing an offset indicating a starting position for the piece of content in the structured data file and a depth value of the piece of content indicating a level in a hierarchy of tags in the structured data file; and

means for formatting said record in a way that allows data to be accessed using both said intermediate file and the structured data file without traversing the entire structured data file to determine the depth value.

27. The apparatus of claim **26**, wherein said means for creating a record comprises:

means for creating a binary mask format (BMF) record in the intermediate file, the BMF record corresponding to said piece of content; and

means for storing an offset in said BMF record indicating a position for said piece of content relative to a point in said structured data file.

28. The apparatus of claim **26**, wherein said intermediate file is a BMF file.

29. The apparatus of claim **27**, wherein said means for creating a record further comprises:

means for storing a depth of said piece of content in said BMF record, said depth indicating a level in a hierarchy of tags in said structured data file.

30. The apparatus of claim **27**, wherein said means for creating a record further comprises:

means for storing a content type of said piece of content in said BMF record, said content type indicating a type of information for said piece of content.

31. The apparatus of claim **27**, wherein said means for creating a record further comprises:

means for storing a length for said piece of content in said BMF record.

32. The apparatus of claim **27**, wherein said offset indicates a starting position for said piece of content relative to a beginning of the structured data file.

33. An apparatus for efficiently processing a structured data file, the structured data file including one or more pieces of content, comprising:

means for creating a binary mask format (BMF) record in a BMF file, said BMF record corresponding to a piece of content;

12

means for storing an offset in said BMF record indicating a starting position for said piece of content relative to a beginning of the structured data file;

means for storing a depth value of said piece of content in said BMF record, said depth indicating a level in a hierarchy of tags in said structured data file;

means for storing a content type of said piece of content in said BMF record, said content type indicating a type of information for said piece of content; and

means for storing a length for said piece of content in said BMF record;

wherein said BMF record may be accessed to determine the depth value without traversing the entire structure data file.

34. The apparatus of claim **33**, further comprising:

means for storing a modification indicator for said piece of content in said BMF record,

said modification indicator indicating if said BMF record has been modified and said modification indicator initially set to indicate that no modification has been made.

35. The apparatus of claim **33**, further comprising:

means for storing an insertion indicator for said piece of content in said BMF record, said insertion indicator indicating if said BMF record has been inserted between two existing BMF records and said insertion indicator initially set to indicate that said BMF record has not been inserted between two existing BMF records.

36. The apparatus of claim **33**, further comprising:

means for storing an end of document indicator for said piece of content in said BMF record, said end of document indicator indicating if said BMF record corresponds to a last piece of content in said structured data file.

37. The apparatus of claim **33**, further comprising:

means for storing a current record in use field for said piece of content in said BMF record, said current record in use field indicating whether said piece of content has been deleted.

38. The apparatus of claim **33**, wherein said content type is a tag, a property, text, or a comment.

39. The apparatus of claim **33**, wherein said content type is a starting tag, property name, property value, text, comment, or processing instruction.

40. The apparatus of claim **33**, wherein said content type is a starting tag, ending tag, property name, property value, text, comment, processing instruction, markup declaration name, markup declaration value, external reference, or property name value pair.

41. The apparatus of claim **33**, wherein said structured data file is an extensible markup language file.

42. A program storage device readable by a machine, tangibly embodying a program of instructions executable by the machine to perform a method for efficiently processing a structured data file, the structured data file including one or more pieces of content, the method comprising:

creating a record in an intermediate file for a piece of content in the structured data file, the record containing an offset indicating a starting position for the piece of content in the structured data file and a depth value of the piece of content indicating a level in a hierarchy of tags in the structured data file; and

formatting said record in a way that allows data to be accessed using both said intermediate file and the structured data file without traversing the entire structured data file to determine the depth value.

US 7,133,857 B1

13

43. A program storage device readable by a machine, tangibly embodying a program of instructions executable by the machine to perform a method for efficiently processing a structured data file, the structured data file including one or more pieces of content, the method comprising:

for each relevant piece of content in the structured data file;
creating a binary mask format (BMF) record in a BMF file, said BMF record corresponding to said piece of content;

storing an offset in said BMF record indicating a starting position for said piece of content relative to a beginning of the structured data file;

5

14

storing a depth value of said piece of content in said BMF record, said depth value indicating a level in a hierarchy of tags in said structured data file;

storing a content type of said piece of content in said BMF record, said content type indicating a type of information for said piece of content; and

storing a length for said piece of content in said BMF record wherein said BMF record may be accessed to determine the depth value without traversing the entire structured data file.

* * * * *